National Aeronautics and Space Administration

Virtex-4VQ Dynamic and Mitigated Single Event Upset Characterization Summary

Gregory Allen
Jet Propulsion Laboratory
Pasadena, CA

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

JPL Publication 09-4 01/09

National Aeronautics and Space Administration

Virtex-4VQ Dynamic and Mitigated Single Event Upset Characterization Summary

NASA Electronic Parts and Packaging (NEPP) Program
Office of Safety and Mission Assurance

Gregory Allen
Jet Propulsion Laboratory
Pasadena, CA

with help from other members of the
Xilinx Radiation Test Consortium

NASA WBS: 724297.40.43
JPL Project Number: 103982
Task Number: 03.04.01

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109

http://nepp.nasa.gov

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration Electronic Parts and Packaging (NEPP) Program.

Table of Contents

1 Introduction

This report is the result of funding by the NASA Electronic Parts and Packaging Program (NEPP) and the combined efforts of members within the Xilinx Radiation Test Consortium (XRTC), sometimes known as the Xilinx Single Event Effects (SEE) Test Consortium. The XRTC is a voluntary association of aerospace entities, including leading aerospace companies, universities and national laboratories, combining resources to characterize reconfigurable, field programmable gate arrays (FPGAs) for aerospace applications. Previous publications of Virtex-4 radiation results are for commercial (non-epitaxial) devices; see, for example, Refs. 1–5. A notable exception is Ref. 6, which presents XRTC upset measurements of storage elements in the PowerPC405s in the XQR4VFX60.

This work represents a continuation of the efforts reported in the "Virtex-4QV Static SEU Characterization Summary" [7]. The contents of this report describe various Single Event Functional Interrupt (SEFI) and Single Event Upset (SEU) modes seen while dynamically exercising the clocked resources within Virtex-4 devices and the corresponding mitigation techniques related to the aforementioned observed SEFI modes.

2 Virtex-4 Overview

The Xilinx Virtex-4 device is a 1.2 V static random access memory (SRAM)-based, in-system, reconfigurable FPGA. The Virtex-4 architecture includes five major, programmable block types optimized for specific functions:

- The Configurable Logic Blocks (CLB) provide functional elements for combinatorial and synchronous logic, including configurable storage elements and cascadable arithmetic functions.
- The Digital Signal Processing (DSP) Slices provide advanced arithmetic and comparison functions, including multiply and accumulate.
- The Block Memory modules provide large 18-kbit storage elements of true dual Port RAM.
- The Digital Clock Manager (DCM) blocks provide clock frequency synthesis and de-skew.
- The bidirectional Input/Output Blocks (IOB) have optional Single Data Rate (SDR) or Double Data Rate (DDR) registers and serializers and deserializers (SERDES) enabling support for many industry input/output (I/O) standards, plus selectable drive strengths and digitally controlled output impedance.

The high reliability, radiation hardened Virtex-4QV product line includes three sub-family (or platform) architectures:

- XQR4VLX200: The LX platform emphasizes logic resources. With more than 200,000 logic cells, the LX200 is optimized for high-density common digital logic applications.
- XQR4VSX55: The SX platform provides a higher ratio of DSP to CLB logic slices, optimizing this architecture for DSP intensive designs.
- XQR4VFX60 and XQR4VFX140: The FX platform includes dual PowerPC 405 processors, optimizing this architecture for embedded processing applications.

Table 1 provides architecture resource values for the aforementioned Virtex-4QV devices.

Table 1. Architecture Resources of the Virtex-4QV Products

	Description	XQR4V SX55	XQR4V FX60	XQR4V FX140	XQR4V LX200
CFG*	Configuration bits (millions)	15.4	14.5	34.5	43.0
BRAM	Block memory bits	5,898,240	4,276,224	10,174,464	6,193,152
LOGIC	Slices (2 lookup tables/slices)	24,576	25,280	63,168	89,088
DSP**	18 × 18 MACs	512	128	192	96
PPC	PowerPC405 processors	-	2	2	-
DCM	Clock managers	12	12	20	12
MGT***	High-speed transceivers	-	N/A	N/A	-
IOBs	Input/output blocks	640	576	896	960

* = Only real memory cells in the Configuration Bit Stream are counted here (not counting BRAM)
** = MAC = multiply-and-accumulate block for digital signal processing (DSP)
*** = MGTs are not supported for Virtex-4QV devices

3 Dynamic and Mitigated Test Overview

Complex devices such as processors, FPGAs, and complex memories often require multiple test iterations at different levels of investigation and operation in order to develop a complete characterization of the device's radiation response. Generally speaking, a dynamic test is a test that requires the element within the device under test (DUT) being characterized to be clocked or functionally exercised. A static cross section of the resources fundamental elements (e.g., control registers) may be developed, but the primary objective is to define the various SEFI modes that the element exhibits. As the Virtex-4 static report provided the static SEE sensitivity of the Virtex-4's fundamental elements, this report aims to provide a comprehensive description of the dynamic failure modes seen in the various resources tested within the device. Once the SEFI modes have been characterized (this characterization may or may not include a cross section versus effective linear energy transfer [LET] curve), a mitigation scheme is developed.

A mitigation scheme may include triplication (triple modular redundancy [TMR]) of the configuration logic and/or primitive, scrubbing, watchdog timers, and/or management circuitry. Details of said mitigation methodologies are beyond the scope of this paper, but can be found in [8]-[13]. TMR'ed mitigated cross sections are developed by selecting an LET that will cause upsets in the device, then varying the flux (and therefore the number of upsets in a given scrub cycle). The system error cross section (number of system errors divided by the fluence) times the flux is then plotted against the flux times the raw bit flip rate at the tested LET. If the TMR

mitigation is correctly implemented, the rate of the uncorrected errors should be proportional to the square of the of the raw bit flip rate [13], [14]. The system error rate can then be calculated from the quadratic fit when the system error rate is measured as a function of the per bit rate. A mathematical proof of this theory can be found in [14], but the concept is somewhat intuitive.

4 Resource Details

FPGA technology increasingly provides designers with more and more configurability and built-in functionality and optimization. This advanced functionality often results in complex SEE responses. This section of the paper defines some of the resources that have been tested in order to under understand their SEE susceptibility and eventually how to mitigate that susceptibility.

4.1 Configurable Logic Blocks (CLBs)

The CLB is the fundamental component of the FPGA that provides function generators, registers, and routing controls. The CLB allows for implementation of macros or other functions. Each CLB is connected to a switch matrix and then to the general the routing matrix. A basic understanding of the CLB and the implementation options surrounding it is imperative in mitigating various SEE responses. Each CLB contains four slices. There are four slices in any given CLB and each slice contains two look-up tables (LUT), two registers (flip-flops [FF] or latches), and carry logic. Two of the four slices in each CLB can be configured into 64 bits of distributed RAM (LUTRAM) or 64 bits of shift registers (SRL16). Each LUT can be configured into an arbitrary, user-defined, 4-input Boolean function. The registers can either be configured as edge-triggered D-type flip-flops or level sensitive latches. The control signals include clock (CLK), clock enable (CE), and set/reset (SR). The control signals of the registers, along with control signals of other primitives in the device, are provided logic levels with weak keeper pull-ups and pull-downs known as half-latches. Half-latches are fixed within the FPGA (i.e., they are not controlled by programmable bits and therefore cannot be scrubbed), but are routed to the correct location. The weak resistive ties can be overridden by any hard routing signal.

SRL16s and LUTRAMs can be implemented through the use of LUT logic. In previous versions of the Virtex devices, bitstream readback and re-configuration were not allowed for configuration columns containing these resources, as those actions would cause corruption of those storage elements. In Virtex-4, readback and reconfiguration can mask out the SRL16 and LUTRAM content through the use of the GLUTMASK bit in the control register. GLUTMASK affects how the memory cells are read back. When disabled, dynamic values will be read back from the SRL16 and LUTRAM primitives; when enabled, all zeros will be read back.

4.2 Block Random Access Memory (BRAM)

In addition to the distributed random access memory (RAM), Virtex-4 devices contain 18 kb block RAM (BRAM). Just as in the case of SRL16s and LUTRAMs, BRAM content could be corrupted if dynamically read back. In addition to triplication, a special BRAM scrubber engine should be implemented separate from the configuration manager. In

addition to the actual BRAM content being sensitive to upset, special BRAM integrity bits can be upset that will cause many hundreds of bit flips at once in the BRAM. The BRAM integrity bit can be scrubbed, which will in turn repair the erroneous BRAM content.

4.3 Digital Clock Managers (DCMs)

Digital clock managers (DCMs) are complete system-level clock managers that enable internal and external clock de-skew, phase adjustment, and frequency synthesis. Clock deskew is accomplished via a delay-locked loop that deskews the DCM's output (CLKFB) with respect to the input clock (CLKIN). The locked output pin (LOCKED) is validated when the DCM outputs exhibit the correct frequency and phase. Frequency synthesis is provided with predefined outputs such as double frequency (CLK2X). The test user can also specify any integer multiplier (M) or divide (D) within a specified range. Similar to frequency synthesis, phase shifting allows both course and fine-grained. Course shifting uses 90, 180, and 270 phases of the CLK0 output. Fine-grained phase shifting allows for both static and dynamic shifting of the CLK. This is done with the granularity of moving the phase positively and negatively by a 1/256 period. Dynamic phase shift (PS), M, and D adjustments are made through the dynamic reconfiguration port (DRP). The DRP allows read/write access to configuration memory within a specific functional block requiring reconfiguration. This memory is directly accessible from the FPGA fabric, and is initialized with the value of the corresponding bit in the bitstream.

4.4 Digital Signal Processing (DSP)

The DSP slice (DSP48) is an 18 × 18 bit two's compliment multiplier. The DSP48 operation is dynamic as the slice can adapt to different functions from clock cycle to clock cycle. A 7-bit operation mode register (OPMODE) controls operation in dynamic mode. The DSP block receives two operators A and B. They are either used as two 18-bit multipliers or as the adder's most significant word (MSW) and least significant word (LSW) respectively. The DSP's adder function requires a second 48-bit operator, C. The result is placed in a 48-bit output register, P. Each DSP slice also contains multiplier (18-bit) and adder (48-bit) carry logic as well as carry, clock, and reset control logic.

4.5 First In First Out (FIFO)

The FIFO primitive has various known issues related to the predefined state flags that can impede even basic functionality. The details related to these issues are found in [15]. It is recommended not to instantiate the FIFO primitive and use BRAM fabric instead.

4.6 Input/Output Blocks (IOB) and Digitally Controlled Impedance (DCI)

Virtex-4 FPGAs provide a series of configurable drivers and receivers (SelectIO), which allow for a variety of single and differential I/O standards. Each I/O tile contains two IOBs, two ILOGIC blocks and two OLOGIC blocks. An IOB consists of an input buffer, tristate output buffer, and connection to the pad. The I/O buffers have direct connections to the ILOGIC/OLOGIC blocks. The ILOGIC blocks consist of four storage elements and a programmable absolute delay element. OLOGIC blocks consist of six storage elements,

three registers for tri-state control, and three for data output. The number of banks available is dependent on the device size and can be found in the datasheet.

DCI control is available to adjust the output impedance or input termination to match transmission line impedance. DCI dynamically adjust the impedance of the I/O to equal the external reference resistance. There are three modes of DCI operation: “as required,” “continuous,” and “freeze,” all of which respond differently in a radiation environment.

4.7 PowerPC and MicroBlaze Processors

Virtex-4 FX devices contain from one to two hardcore PowerPC processors, and all devices are capable of being implemented by the soft-core Microblaze processor IP. Processor mitigation has been intentionally left out of this report due to the sheer complexity and in-depth description necessary to describe the application, testing, and mitigation techniques. Details concerning the PowerPC and MicroBlaze can be found in [16] and [17] respectively.

5 Experimental Setup

Figure 1 shows the test setup in vacuum. Details of the test board used for all experiments can be found in [18]. When in vacuum five, 40-pin bulkheads were used to run five of the six communication cables through the vacuum chamber. The sixth was run through the 50-pin D-Sub connector provided by Texas A&M University (TAM). Three parallel cables were also sent through the 50-pin connectors (one for a DUT readback Parallel-IV cable, one for a motherboard/DUT design programming Parallel-IV cable, and one for the temperature monitoring circuit). A mounting platform with integrated power breakout cables was used for mounting the motherboard to the rotating chassis in the vacuum chamber, and for extracting the four power supplies from the 40-pin cable. The four supplies were sent through the vacuum chamber bulkhead over Bayonet Neill Concelman BNC connectors, then re-integrated to the 40-pin cable. Force and sense were tied together at the power supply (HP6629) for all four supplies, and provided the necessary 2.5 V, 3.3 V, and 3.3 V I/O for the motherboard; the last supply was used to control heater strips attached to the back of the daughter card. The receiver/driver cards were powered by the 3.3 V of the motherboard I/O. The 5 V for the Parallel-IV cables and temperature sensor circuit were powered by an external Agilent E3610A, and also run through a BNC bulkhead (provided by TAM). Typically, the DUT power supply was an HP6623, which provided three supplies, with currents of 5 A, 10 A, and 2 A on supplies one, two, and three respectively. Supply one provided 2.5 V to the auxiliary voltage supply VAUX while supply two provided 1.2 V to VINT, and supply three provided 3.3 V to two the I/O voltage VCCO DUT I/O banks that talk to the motherboard. Force and sense were tied together at the bulkheads on supplies one and three (which were run through BNC bulkhead feedthroughs). High current cables were used to run force for supply two in through a 40-pin cable and bulkhead connector, and were separated back into banana cables with a second custom 40-pin connected to an additional banana cable (20 pins were used for power and 20 for ground). Sense for supply two was sent through a BNC over banana cables and connected to force at the daughter card.

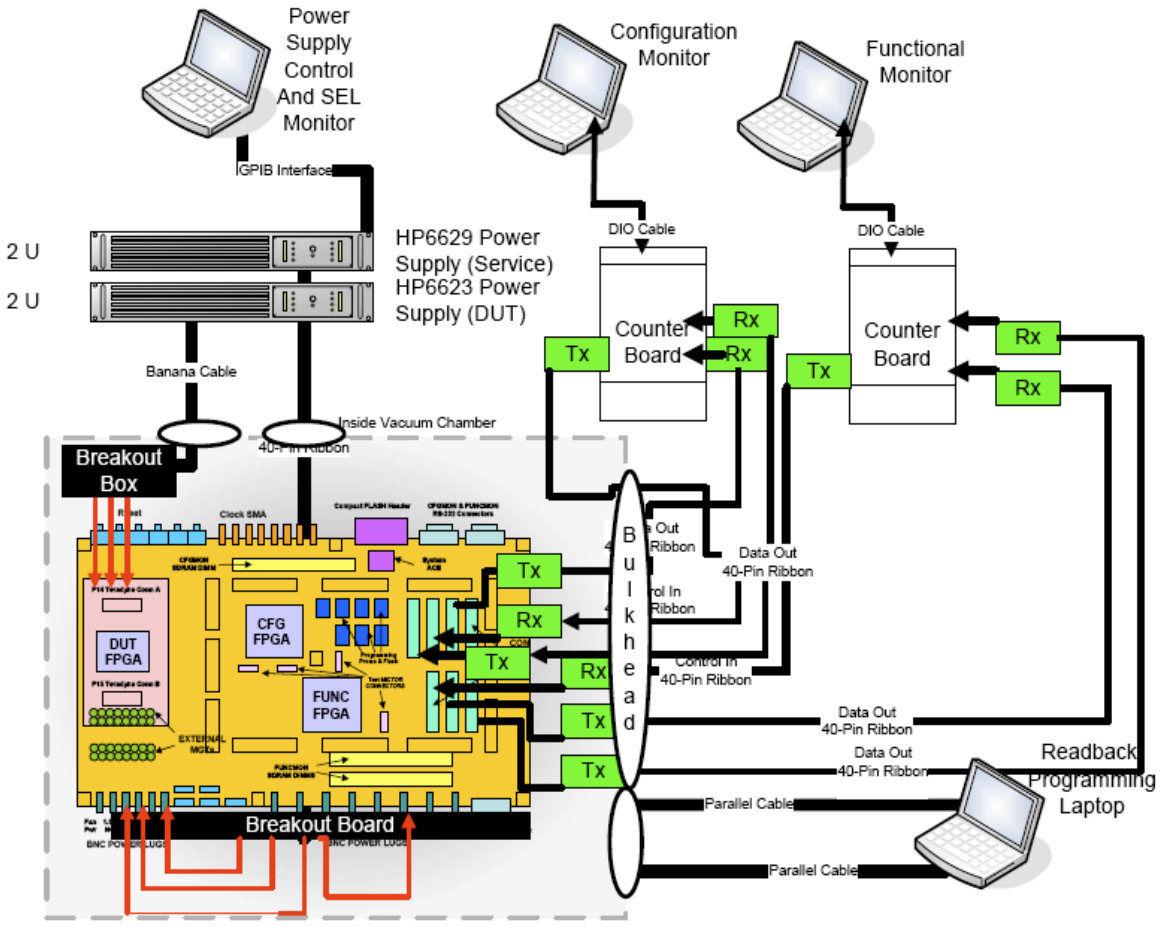

Figure 1. Experimental setup in vacuum at Texas A&M. Everything outside of the vacuum chamber border is setup in the SEE user room.

The setup for in-air testing was essentially the same as in vacuum, the main exception being that the adapted connections for getting through the bulkheads were discarded. Also, USB programming cables were used via high-speed hubs for the in-air irradiations.

6 Dynamic and Mitigated Test Results

6.1 DCM

Dynamic and mitigated tests were performed at Texas A&M's cyclotron in February and August, 2008. The initial dynamic DCM test took place using the XQRV4SX55. Table 2 shows the beams used for the test. The design consisted of all eight of the device's DCMs being instantiated and monitored. Although the primary areas of focus were the CLKFX and LOCK output pins of the DCM, CLKDV and CLK0 were monitored via the FuncMon FPGA. The first attempt of dynamic testing was implemented with GLUTMASK both enabled and disabled. With GLUTMASK enabled (set to '0') scrubbing skipped over SRL16s and LUTRAMs, but also had no access to the DRP bits. It this case, when a DCM failure occurred, the M, D, and PS values were manually re-written via DRP, a reset DCM

was asserted, and the failure was cleared. This cross section is shown in Figure 2. In the second design, GLUTMASK was disabled (set to '1'). In this case, SRL16s and LUTRAMs had to be removed from the design, as scrubbing the configuration would corrupt those elements with GLUTMASK disabled. But, with GLUTMASK disabled, the scrubber was able to scrub the DCM bits related to DRP, which were previously inaccessible with GLUTMASK enabled. In this case, all DCM failures recovered with the only intervention being the scrubbing engine (i.e., all DCM upsets were transient). Another important lesson learned in dynamically testing the DCMs was that the LOCKED pin is based only on the CLK0 pin (i.e., not the CLKDV, CLK90, CLK180, etc.). This becomes important when implementing a DCM mitigated design.

Table 2. Ions used for dynamic DCM testing. Effective LETs were acquired using either degrader or angle of incidence.

Ion (25 MeV/u)	Effective LET (MeV-cm²/mg)	Effective Range (μm Si)*
Ne	5.3	160.7
Ar	8	211.5
Ne	9	47
Ar	9.9	126.1
Ar	13.9	92.4
Ar	16.7	66.3
Ar	22	38.2
Kr	33.1	77.1

*** The effective range is the range after the predefined layer file.**

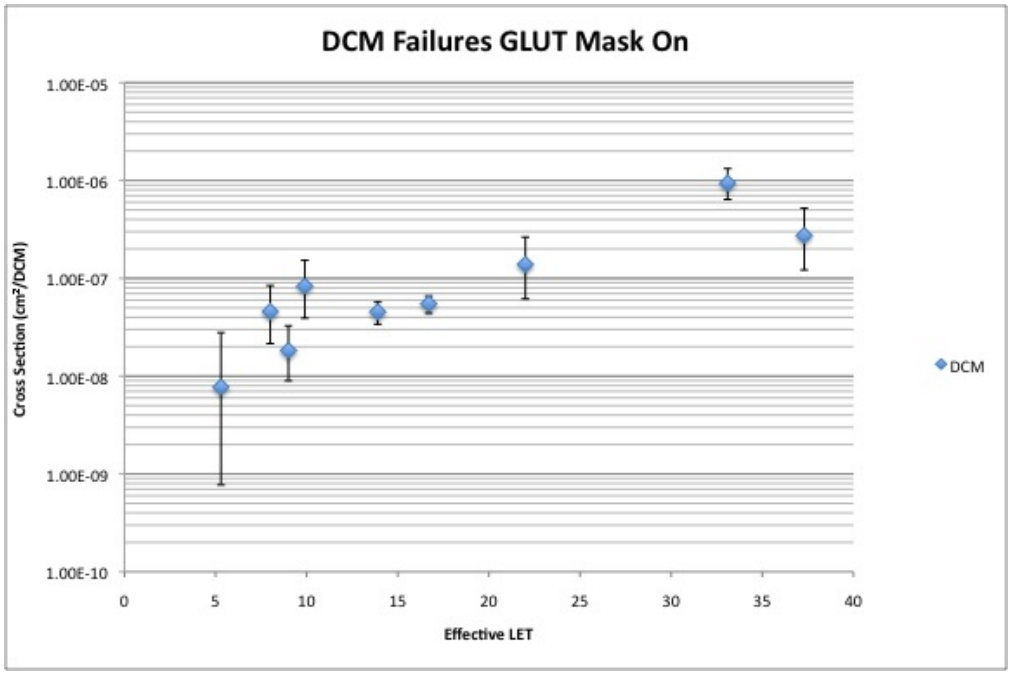

Figure 2. DCM failures occurring with GLUTMASK enabled. Either DCM reset, rewriting the M or D values through DRP, or disabling GLUTMASK and scrubbing fixed all of the failures.

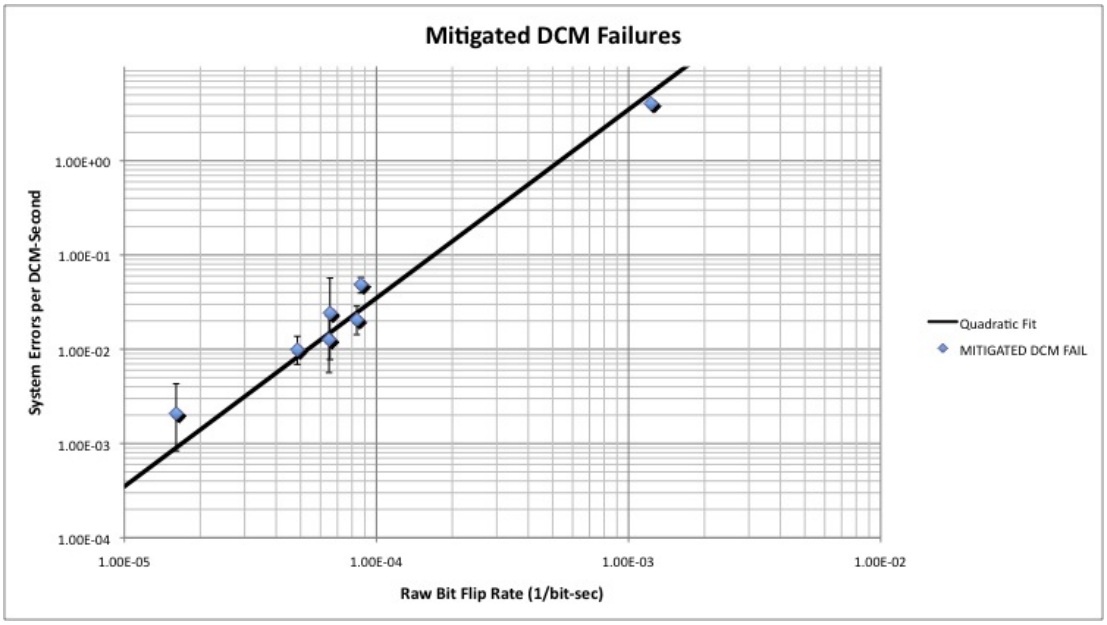

Figure 3. Shows system error rate per DCM versus the raw bit flip rate for the mitigated DCM design.

The mitigated design consisted of three separate DCMs, each in their own clock domain with no clock voting taking place (although control logic was triplicated). The design was implemented with GLUTMASK disabled (i.e., the DCM's, DRP associated bits were being scrubbed). Each output clock of the DCM fed a 2-bit counter. If a discrepancy occurred in the output counters, the DCM was reset. Throughout the testing, no persistent errors were discovered (i.e., scrubbing repaired any failures). A system failure was counted if two or more DCMs needed to be reset at the same time, where "same time" is defined as happening within a single scrub cycle. It should be noted that prior to beam testing, fault injection was performed on the design to ensure that there were virtually no single points of failure (excluding injected faults that caused device SEFIs). The testing was performed at Lawrence Berkeley National Laboratory's (LBNL) 88-inch cyclotron in vacuum. A 16 MeV per nucleon Argon beam was used providing an LET of 7.27 MeV-cm^2/mg. The range of this beam is 256 um in silicon. Figure 3 shows the system error rate per DCM versus raw bit flip rate for the tested mitigated design. Note the quadratic fit equation:

$$\text{System Error Rate} = 3.5 \times 10^{6} * (\text{Upset Rate})^{2} \quad (1)$$

where the system errors per DCM implies a mitigated DCM circuit (i.e., three DCMs).

The fit shown in Figure 3 shows a quadratic relationship between the system error rate and the raw bit flip rate, as we expect to see [19]. This reinforces the belief that the design is correctly TMRed (as implied from the fault injection testing). A space rate for this design can now be extrapolated by extending this trend down to lower flux levels encountered in space. For example, in a geosynchronous GEO orbit, 4.28 bit flips per device-day are expected for an SX55 [7]. This equates to 3.2×10^{-12} bit flips per bit-second. Extrapolating the quadratic fit from equation 1 to the expected bit flip rate, we get a failure rate of approximately 1×10^{-2} system errors per DCM circuit-century (i.e., once every 10,000

years). This rate compared to the device SEFI rate, which for a GEO orbit is approximately one per century, implies a more than satisfactory rate for this mitigation strategy.

6.2 DSP

Pseudo-static testing of the DSP blocks took place in October 2008 at Texas A&M's cyclotron. Table 3 shows the ions and corresponding LETs and ranges. Testing was performed on the Virtex-4 SX55; the tests utilized all of the device's 512 DSP slices. The test methodology was as follows: reset and reload the DSPs, execute a predefined operation and check the result, irradiate to a preset fluence, clock out the operational and resultant registers, and count errors. If a device SEFI occurred during irradiation, that run was discarded. Two types of events were observed: typical register-bit SEU, and "clear events" where the register was zeroed out. Figures 4–6 show the M, C, and P register clear events, the A and B register upsets, and the M-C-P register upsets respectively.

Table 3. Ions used for dynamic DSP testing.

Ion (25 MeV/u)	Effective LET (MeV-cm²/mg)	Range (µm Si)*
Ne	1.9	799
Ne	3.6	799
Ar	6.2	493
Kr	22.7	332
Kr	34.9	332
Kr	35.3	332
Xe	46.3	286
Xe	49.3	286

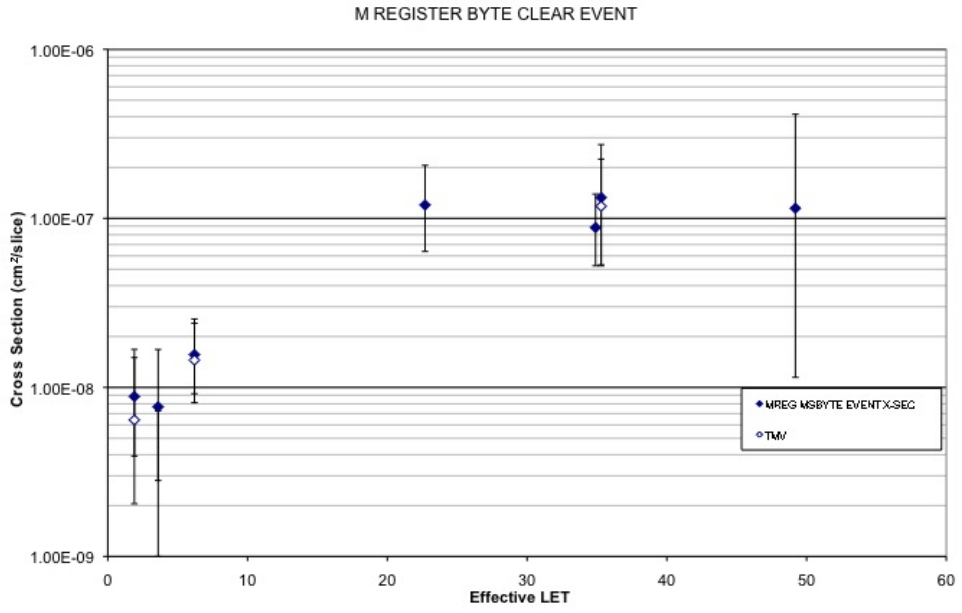

Figure 4. Shows the M register clear event (zero reset) cross section per DSP slice versus effective LET. Two sets of data points are shown. The solid black points represent a fully unmitigated design; the white points represent triplicated logic surrounding unmitigated DSP slices.

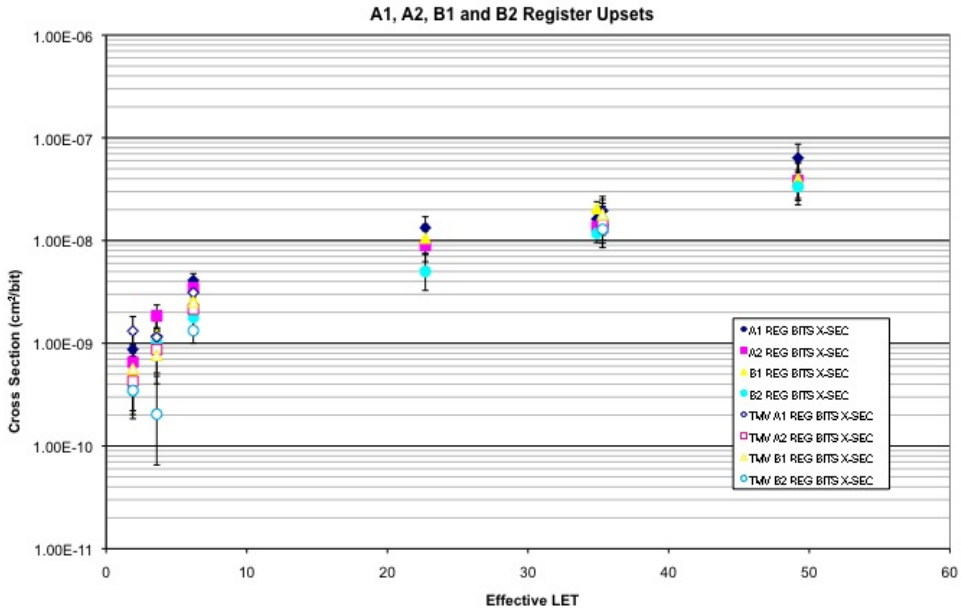

Figure 5. Shows the A and B register per bit upset cross section versus effective LET. Two sets of data points are shown. The solid points represent a fully unmitigated design; the outlined points represent triplicated logic surrounding unmitigated DSP slices.

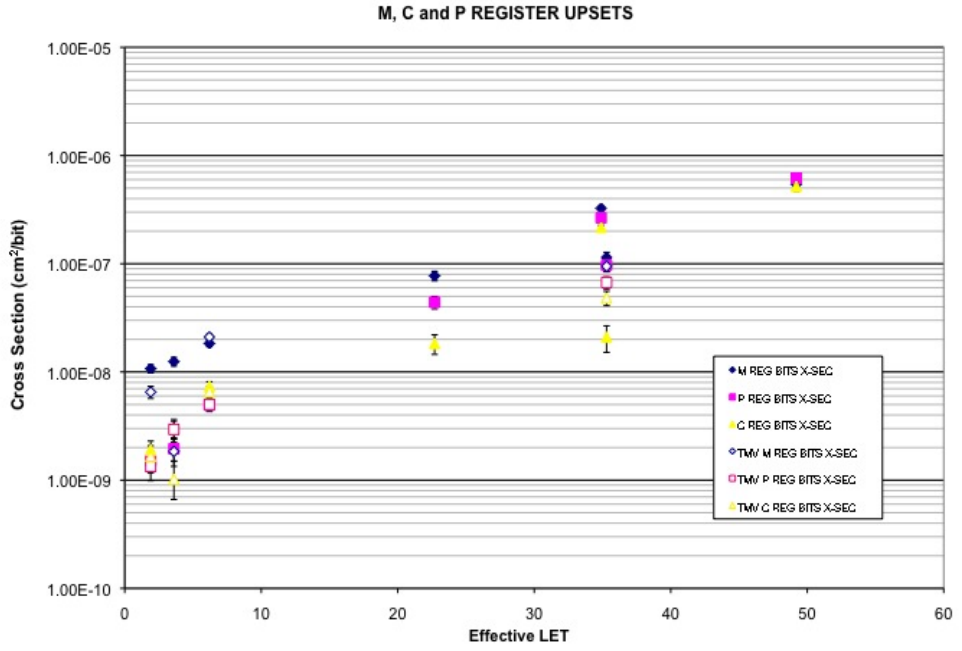

Figure 6. Shows the M, C, and P register per bit upset cross section versus effective LET. Two sets of data points are shown. The solid points represent a fully unmitigated design; the outlined points represent triplicated logic surrounding unmitigated DSP slices.

6.3 IOB and DCI

The main purpose of the dynamic IOB testing was not to characterize the failure rate of the I/O (as that can be estimated with configuration bit upset rates), but to discover any SEFI mode that may exist that affects the I/O of the device. A specific SEFI of interest that was seen in Virtex-II Pro caused bank-wide failures of I/O, whereby every pin on a given I/O bank would tristate until the contributory upset(s) were scrubbed. The dynamic design was similar to that found in [14]. The dynamic test consisted of routing a 33 MHz clock signal (with a 50% duty cycle) via the LVCMOS25 standard from the Functional Monitor FPGA through a DUT input, through a minimal amount of routing, and back to the Functional Monitor for comparison. Forty-seven I/O set pairs were used, with careful diligence put into their assignment to ensure minimal DUT routing and matched I/O banks. Due to mismatched routing lengths in the motherboard of the single-ended I/O, an oversampling clock was used to check for a signal mismatch (i.e., after the edge of the data clock, the over sample clock counted up to a preset number before declaring a mismatch). There were three sets of counters implemented: the first was a simple mismatch counter, counting the total number of mismatches seen by the Functional Monitor; the second was a stream error counter, which incremented by one every time there was a set of mismatch errors greater than 8. The final counter was a transient error counter, which was intended to count transient errors. The transient detector consisted of a pair of edge detectors (positive and negative edge flip flops) and a small state machine that checked for a set of edges and the expected data state. The mitigated (triplicated) test consisted of the same functional monitor configuration with the DUT design run through the TMRTool flow.

As previously mentioned, there are three DCI modes available in the Xilinx tool flow: “as required,” “quiet,” and “continuous.” The aforementioned IOB test was performed with each of the three DCI modes separately activated. A significant increase was seen in the bank-wide failures with each of the DCI modes. As such, it is recommended not to use DCI in one’s design. If DCI is necessary for proper operation of a design, “as required,” the most robust of the three, is the suggested mode of operation.

6.4 CLB

A test dubbed the “CLB test” was developed as a proof of concept of a full triplication and configuration management mitigation scheme. Because the test is designed to test mitigation concepts, no dynamic test was necessary. The design was implemented with the intent to mimic a complex state machine with I/O variables. It was also designed with the intent of filling the device in order to increase the error cross section and shorten the length of time required to acquire test data. There are 16 counter modules within the design. The error counter for each module is displayed and recorded in 16 of the functional monitor’s output registers. Each module contains 32, 16-bit counters. Each counter increments by a larger number incrementally (e.g., 1, 2, 3…). The output of the module is a 32 × 1 × 16 multiplexor. The output of the top level is a 16 × 1 × 16 multiplexor with a final register stage. The design is fully triplicated, including I/O pins. For each data snap shot, all counters are sent four clock pulses. This is to advance new data through the multiplexor pipeline. Then without further clocking the DUT, the design cycles through each counter for each module. The number of word errors for each generator is added together for a total

error count for that module. One system error is counted for each counter that is outputting an incorrect value regardless of how many bits are incorrect.

Figure 7 shows the results of a typical tri-flux test. It should be noted that this example design was not full triplicated correctly. This example was intentionally selected as a discussion point in order to show the strength of the tri-flux test and how to interpret the results. A 25 MeV/u Kr beam was used providing an LET of 22.7 MeV-cm^2/mg and range of 332 um (Si). The test was executed at three fluxes averaging 1E3, 1E4, and 1E5 ions/(cm^2-sec). The red curve indicates a quadratic slope ($y = bx^2$), the relationship we expect to see for a fully triplicated design. The green curve indicates the unit slope ($y = ax$), what we would expect to see if the design had no triplication. The purple curve, and best fit is a summation of the two curves ($y = ax+bx^2$). The constant 'a' in the summation fit indicates the number of untriplicated bits (i.e., single points of failure in the design). In this case, testing indicates that there are approximately 60 single points of failure in the design. Equation 2 shows the fit for the CLB data.

$$\text{System Error Rate} = 60*(\text{Upset Rate}) + 3.5 \times 10^5 * (\text{Upset Rate})^2 \qquad (2)$$

Extrapolating to the raw bit flip rate in a GEO orbit, we calculate a system error rate of approximately 0.6 system errors per device-century. For all intensive purposes, this is equivalent to the SEFI rate of approximately 1 per device-century; this may or may not be acceptable for the given mission.

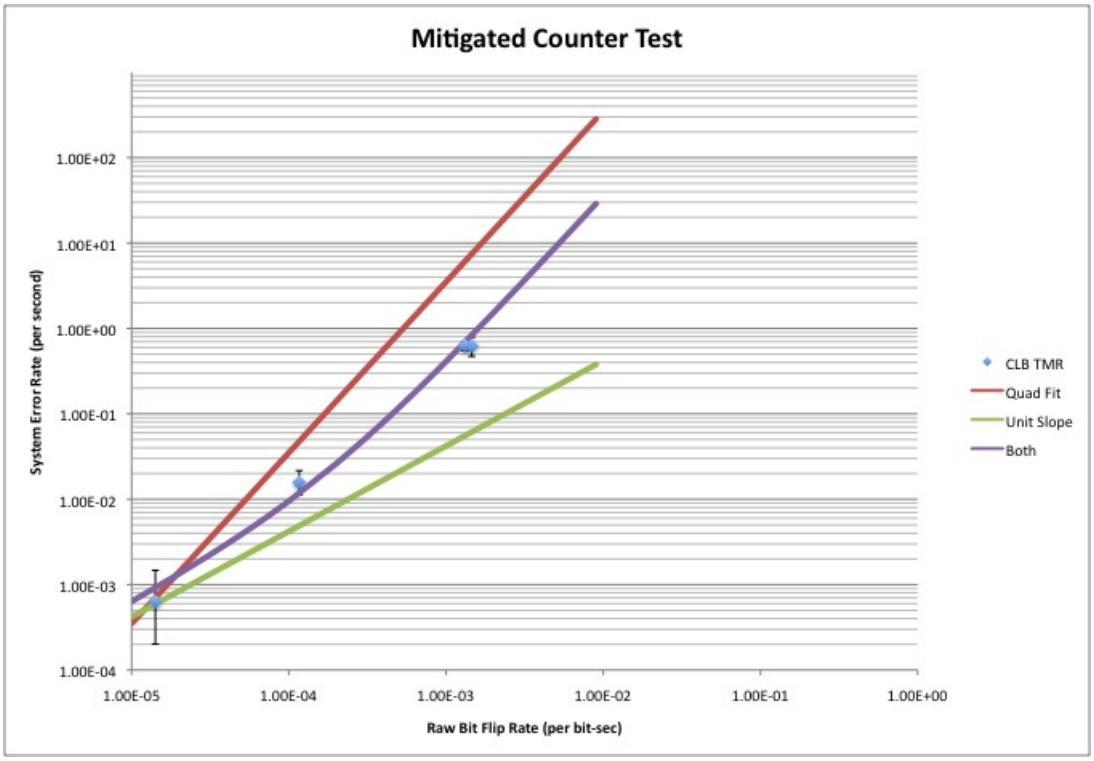

Figure 7. Shows the system error rate versus the raw bit flip rate for the mitigated counter test.

Because fault injection was not performed prior to ion beam testing, costly beam testing was used to show single points of failure. At this point, knowing there are single points of failure, fault injection would be performed to locate those single points of failure and design iterations performed until those points were reduced to a number that provides the application with an acceptable level of risk. Ion beam testing could then be repeated depending on the level of design mitigation verification required by the project or application.

6.5 Embedded Processors

Results for pseudo-static/dynamic testing of the PowerPC can be found in [5]. Various application specific studies have been performed on the microblaze and PowerPC processors, but are too specific to be mentioned here. Generic MicroBlaze testing is scheduled for 2009.

6.6 Half Latches

Dynamic testing of half latches in Virtex-4 took place at Texas A&M in March 2007. Two forums of mitigation for half latch upsets exist for Virtex-4: the first is to remove the half latches with Xilinx's TMRTool; the second is to replace any half latch with a dedicated power post; defining the primitives and setting the environment variables seen below will allow use of dedicated power posts. Half latch upsets were not found to be persistent in Virtex-4.

Environment variables to be set:

XIL_MAP_RETAIN_CONSTANT_FF_CTRL

XIL_PAR_VCC_HARD1_ONLY

Primitive definitions:

```
component VCC
   port
   (
         P : out std_ulogic := '1'
   );
end component;

component GND
   port
   (
         G : out std_ulogic := '0'
   );
end component;
```

7 Conclusions and Future Work

7.1 Conclusions

The Virtex-4QV family of devices provides a great deal of reprogrammable processing power. The reconfigurable cells that provide these devices with their processing power introduce a significant level of risk when applied within a radiation environment due to their SEU susceptibility. Moreover, the integrated primitives allow for a great deal of fast functionality, and also produce complex radiation responses (SEFI modes) that need to be taken into consideration when designing. That being said, the burden of mitigation will lie upon the designer.

The dynamic testing present in this document has shown a great deal of work in providing designers with all of the observed SEE responses seen in our study of the basic clocked components of the Virtex-4QV devices. The reality is, the dynamic test matrix is as infinite as the infinite configurability of the FPGA. The responsibility of developing a mitigation strategy is convoluted by the need to verify that mitigation strategy. Fault injection can go a long way in that verification process, but absolute verification can only be obtained via beam testing.

7.2 Future Work

This document exists as a living document. As more knowledge is gathered and SEE modes observed, the document will be augmented. Work remains to analyze and provide rates for some of the existing data such as half-latches and BRAM group errors. Mitigation strategies are presently being developed and tested, such as DSP mitigation. Processor testing including the MicroBlaze and PowerPC are planned for 2009, but may be included in a separate document. In addition to reporting on radiation data, work is being done to provide designers with mitigation strategies (e.g., what level of mitigation is appropriate) along with mitigation verification techniques (including what level of mitigation verification is appropriate).

8 References

[1] J. George et al., "Single Event Upsets in Xilinx Virtex-4 FPGA Devices," *2006 IEEE Radiation Effects Data Workshop Record*, pp. 109–114, 2006.

[2] H. Quinn et al., "Domain Crossing Errors: Limitations on Single Device Triple-Modular Redundancy Circuits in Xilinx FPGAs," *IEEE Trans. Nucl. Sci.,* vol. 54, no. 6, pp. 2007, Dec. 2007.

[3] M. Berg et al., "Effectiveness of Internal vs. External SEU Scrubbing Mitigation Strategies in a Xilinx FPGA: Design, Test, and Analysis," presented at RADECS 2007, Sept. 2007.

[4] A. Vera et al., "Dose Rate Upset Investigation on the Xilinx Virtex IV Field Programmable Gate Arrays," *2007 IEEE Radiation Effects Data Workshop Record*, pp. 172–176, 2007.

[5] G. Allen et al., "Upset Characterization and Test Methodology of the PowerPC405 Hard-Core Processor Embedded in Xilinx Field Programmable Gate Arrays," *2007 IEEE Radiation Effects Data Workshop Record*, pp. 167–171, 2007.

[6] D. M. Hiemstra, F. Chayab, and Z. Mohammed, "Single Event Upset Characterization of the Virtex-4 Field Programmable Gate Array Using Proton Irradiation," *2006 IEEE Radiation Effects Data Workshop Record*, pp. 105–108.

[7] G. R. Allen, G. Swift, and C. Carmichael, "Virtex-4QV Static SEU Characterization Summary," http://parts.jpl.nasa.gov/docs/NEPP07/NEPP07FPGAv4Static.pdf
[8] G. Miller, C. Carmichael, and G. Swift, "Single-Event Upset Mitigation for Xilinx FPGA Block Memories," http://www.xilinx.com/support/documentation/application_notes/xapp962.pdf
[9] B. Bridgford, C. Carmichael, C. Tseng, "Single-Event Upset Mitigation Selection Guide," http://www.xilinx.com/support/documentation/application_notes/xapp987.pdf
[10] C. Carmichael, and C. Tseng, "Correcting Single-Event Upsets in Virtex-4 Platform FPGA Configuration Memory," http://www.xilinx.com/support/documentation/application_notes/xapp988.pdf
[11] C. Carmichael, and C. Tseng, "Correcting Single-Event Upsets with a Self-Hosting Configuration Management Core," http://www.xilinx.com/support/documentation/application_notes/xapp989.pdf
[12] G. Miller, C. Carmichael, and G. Swift, "Single-Event Upset Mitigation Design Flow for Xilinx FPGA PowerPC Systems," http://www.xilinx.com/support/documentation/application_notes/xapp1004.pdf
[13] P. Adell, and G. Allen, "Assessing and Mitigating Radiation Effects in Xilinx FPGAs," http://parts/docs/NEPP07/NEPP07FPGARadiationEffectsGuideline.pdf
[14] G. Swift, S. Rezgui, C. Carmichael, et al. "Dynamic Testing of Xilinx Virtex-II Field Programmable Gate Array (FPGA) Input Output Blocks (IOBs)." http://parts.jpl.nasa.gov/docs/NSREC04/NSREC04_E3.pdf
[15] "AR #22462—Virtex-4—Why are the FIFO16 flags not working correctly?" http://www.xilinx.com/support/answers/22462.htm
[16] PowerPC Processor Reference Guide. January 19, 2007. http://www.xilinx.com/support/documentation/user_guides/ug011.pdf
[17] MicroBlaze Processor Reference Guide. http://www.xilinx.com/support/documentation/sw_manuals/mb_ref_guide.pdf
[18] Document within private email from Sam Minger of Seakr Engineering. Sent July 1, 2005.
[19] L. Edmonds, "Analysis of SEU Rates in TMR Devices." Internal Document. November 22, 2008.